

# MapSQ: A MapReduce-based Framework for SPARQL Queries on GPU

Jiaying Feng<sup>1,3</sup> Xiaowang Zhang<sup>1,3</sup> Zhiyong Feng<sup>2,3</sup>

<sup>1</sup>School of Computer Science and Technology, Tianjin University, Tianjin 300350, P.R.China

<sup>2</sup>School of Computer Software, Tianjin University, Tianjin 300350, P. R. China

<sup>3</sup>Tianjin Key Laboratory of Cognitive Computing and Application, Tianjin 300350, P.R.China

## ABSTRACT

In this paper, we present a MapReduce-based framework for evaluating SPARQL queries on GPU (named MapSQ) to large-scale RDF datasets efficiently by applying both high performance. Firstly, we develop a MapReduce-based Join algorithm to handle SPARQL queries in a parallel way. Secondly, we present a coprocessing strategy to manage the process of evaluating queries where CPU is used to assigns subqueries and GPU is used to compute the join of subqueries. Finally, we implement our proposed framework and evaluate our proposal by comparing with two popular and latest SPARQL query engines gStore and gStoreD on the LUBM benchmark. The experiments demonstrate that our proposal MapSQ is highly efficient and effective (up to 50% speedup).

## Keywords

RDF; SPARQL; Query; MapReduce; GPU

## 1. INTRODUCTION

SPARQL is the standard query language over RDF datasets. With the rapid growth of RDF data and high complexity of SPARQL, processing RDF data effectively becomes a significant challenge. For instance, a SPARQL query as follow:

*Q: SELECT ?person WHERE { ?person hasJob ?job. ?job workAt "Hospital". }*

Many approaches have been proposed to answer SPARQL queries, such as centralized engine gStore[3], and distributed engine[2]. Their efficiency depends on the calculative capabilities of CPU. However, the performance of CPU has almost reached its peak recently.

In recent years, graphic processing units (GPUs) has been adopted to process large-scale data, and their highly parallel structure makes them more efficient than general-purpose CPUs [4]. [1] develops a MapReduce framework on GPU. MapReduce is a parallel programming model, which excels at large-scale data processing. Combining the characters of MapReduce and RDF datasets, we develop MapReduce-base join algorithm to improve the efficiency of RDF data processing.

SPARQL queries contains a series of triple patterns. Consider the following two triple patterns  $P_1(?person, hasJob, ?job)$  and

$P_2(?job, workAt, "Hospital")$ . Each triple pattern maps a partial matches set, shown as Table 1(a) and 1(b). Both of them hold key and value. The key of them is their shared variable  $?job$ , and the value is other variables. Final results can be obtained by joining partial results, shown as Table 1(c).

**Table 1: Example of MapReduce-based Join**

(a) $Tp_1$		(b) $Tp_2$	
Key1	Value1	Key2	Value2
Proffesor	Anny	Doctor	"Hospital"
Doctor	Jim	Nurse	"Hospital"
Nurse	Susan		

(c) $RS$	
Key	Value
Doctor	Jim, "Hospital"
Nurse	Susan, "Hospital"

In this paper, we propose a SPARQL queries framework on GPU called MapSQ(MapReduce-based framework for SPARQL Queries), which implements results join based on MapReduce. Our work has three major contributions: (1) we propose a MapReduce-based join algorithm on GPU to match intermediate results, (2) we develop the MapSQ framework for improving SPARQL queries performance, and (3) extensive experiments confirm the efficiency and effectiveness of our method.

## 2. MAPSQ FRAMEWORK

**MapReduce-based Join:** Join operation is the basic function in database system and has been researched extensively in CPU-based architecture. Many approaches have been designed by plain join algorithms, such as nested-loop join. We develop a novel method to improve the performance of centralized RDF engine.

In our work, join operation merges partial results after each triple patterns is matched. The output of matching is delivered to join if triple patterns contain one shared variables.

Algorithm 1 shows the MapReduce-based join on GPU. In order to make full use of parallelization advantages, the join operation divided into three phases: Map, Sort and Reduce duplicate. When mapping phase, we set flag according to the shared variable's position in triples and split each triple. The flag, including left or right, contributes to reduce unnecessary computation in the stage of reducing duplicate. GPU's Single Instruction Multiple Data architectures contribute to accelerate cartesian product in parallel.

**MapSQ Framework:** We present design and implementation of MapSQ framework, shown in Figure 1. The bottom of framework is GPU and CPU-based gStore, which is used as a black box for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

XXX XXX

© 2016 ACM. ISBN 0-12345-67-8/90/01...\$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

**Algorithm 1** MapReduce-based Join

**Input:**  $T_{p1}$  and  $T_{p2}$  are the partial matches of each triple pattern.

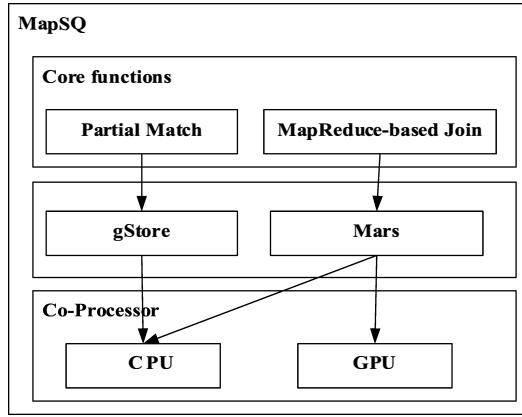
**Output:**  $RS$ , the join results of two triple patterns.

```

1: function Map( $T_{p1}, T_{p2}$ )
2:   for each item  $\in (T_{p1}, T_{p2})$  do
3:     split and emit intermediate.
4:   end for
5:   sort intermediates.
6:   function ReduceDuplicate
7:   for each key do
8:     if value1's label is RIGHT and value2's label is LEFT then
9:       put the key-value pair into  $RS$ .
10:    end if
11:  end for
12: return  $RS$ ;

```

answering SPARQL queries.



**Figure 1: MapSQ Framework**

When answering SPARQL queries, there are two steps to process queries: partial matching and MapReduce-based join. The first step is each triple pattern matches the partial results through centralized RDF engine in parallel, which adopts gStore to dispose queries. The second step is computing the final results by assembling partial matchings. MapReduce-based join Algorithm is implemented on GPU.

### 3. EXPERIMENTS AND EVALUATIONS

In this section, we present the performance of MapSQ framework in comparison with CPU-based methods, including gStore and gStoreD. The purpose of the experiments is to evaluate the improvement of efficiency. Our experiments were carried out on a PC with a GTX590 GPU and an Intel(R) Core(TM) i7-2600 Quad processor running Linux ubuntu14.04. The CPU main memory is 2GB, and the GPU device memory is 1536MB.


We used LUBM<sup>1</sup> as the benchmark datasets in our experiments. LUBM is an ontology for university domain and generates arbitrary scale dataset. We choose 5 benchmark queries to test queries performance. gStore and gStoreD are better performance methods in centralized RDF engine.

Table2 and Figure 2 show the query response time of gStore, gStoreD and MapSQ. And the speedup is computed among them. Our MapReduce-based join algorithm is faster than the join operation of gStore and gStoreD. Comparing with gStore and gStoreD,

<sup>1</sup><http://swat.cse.lehigh.edu/projects/lubm/>

**Table 2: join time(ms)**

Query	gStore	gStoreD	MapQS	SpeedUp <sub>g</sub>	SpeedUp <sub>D</sub>
Q1	1677	2079	1429	1.17	1.45
Q2	3886	4131	3081	1.26	1.34
Q3	433	442	351	1.23	1.25
Q4	635	698	339	1.87	2.05
Q5	7254	10875	6292	1.15	1.72



(a) Queries Response Time

(b) SpeedUp

**Figure 2: Query performance**

MapSQ queries response time reduce 15%-50%, shown in Figure 2(a). the speedup among them shown in Figure 2(b). Experiment results show that our framework performs better than the algorithms on CPU and can efficiently speed up SPARQL queries answering.

Experiment results show that our framework performs better than the algorithms on CPU and can efficiently speed up SPARQL queries answering. Furthermore, there has been a significant increase in different queries. For instance, we observe  $Q_4$ , which reduce to the half of CPU-based engine. In this scene, MapSQ shortens query response time especially large dataset scale.

### 4. CONCLUSIONS

In this paper, we proposed a MapReduce-based framework for SPARQL queries on GPU and the preliminary experiments show our proposal can evaluate SPAQRL queries on large-scale RDF datasets effectively and efficiently (up to 50%, without further optimization). We think that our proposal will bring new methods in processing large-scale RDF data by integrating parallel architecture with new hardware. This work is supported by the programs of the National Key Research and Development Program of China (2016YFB1000603) and the National Natural Science Foundation of China (61672377).

### 5. REFERENCES

- [1] B. He, W. Fang, Q. Luo, N. Govindaraju, and T. Wang. Mars: a MapReduce framework on graphics processors In: *Proc. of PACT'08*, pp. 260-269, 2008.
- [2] P. Peng, L. zou, MT. Zsu, L. Chen, and D. Zhou. Processing SPARQL queries over distributed RDF graphs. *J. VLDB.*, 25(2): 243-268, 2016.
- [3] L. Zou, J. Mo, L. Chen, MT. Zsu and D. Zhou. gStore: Answering SPARQL queries via subgraph matching. In: *Proc. of VLDB Endowment*, pp.482-493, 2011.
- [4] J. Zhang, B. Zhang, X. Zhang, and Z. Feng. IRSMG: Accelerating Inexact RDF Subgraph Matching on the GPU In: *Proc. of ISWC Poster*, 2016.
- [5] B. He, K. Yang, R. Fang, M. Lu N. Govindaraju, Q. Luo, and P.Sander. Relational joins on graphics processors In: *Proc. of SIGMOD'08*, pp. 511-524, 2008.